

# Off-Path Hacking:

## *The Illusion of Challenge-Response Authentication*

Yossi Gilad<sup>\*</sup> and Amir Herzberg<sup>†</sup> and Haya Shulman<sup>‡</sup>

Computer Science Department  
Bar Ilan University  
Ramat Gan, Israel

### ABSTRACT

Everyone is concerned about the Internet security, yet most traffic is not cryptographically protected. The usual justification is that most attackers are only *off-path* and cannot intercept traffic; hence, challenge-response mechanisms suffice to ensure authenticity. Usually, the challenges re-use existing ‘unpredictable’ header fields to protect widely-deployed protocols such as TCP and DNS.

We argue that this practice may often only give an *illusion of security*. We present recent off-path TCP injection and DNS poisoning attacks, enabling attackers to circumvent existing challenge-response defenses. Both TCP and DNS attacks are non-trivial, yet very efficient and practical. The attacks foil widely deployed security mechanisms, such as the Same Origin Policy, and allow a wide range of exploits, e.g., long-term caching of malicious objects and scripts.

We hope that this article will motivate adoption of cryptographic mechanisms such as SSL/TLS, IPsec and DNSSEC, and of correct, secure challenge-response mechanisms.

### 1. INTRODUCTION

Since 1989 [4], experts have been arguing that Internet security requires cryptographic protocols, ensuring security against *Monster-in-the-Middle (MitM) attackers*. A MitM attacker is located on the path of the communicating parties, and can manipulate the communication between them in any way, i.e., intercept, modify, block and inject spoofed packets; see the *MitM Cookie Monster* in Figure 1.

The information security community invested huge efforts in developing cryptographic schemes and protocols, standards and products, providing security against MitM attackers, such as IPsec, SSL/TLS, Secure-BGP and DNSSEC. In spite of all these efforts, and although Internet security is




Figure 1: Off-Path Attacker Network Model.

well recognised to be critical, most Internet traffic is still not cryptographically protected. For example, we found that only about 6% of the TCP traffic is cryptographically protected with SSL/TLS (based on CAIDA dataset of 3 million packets [7]); and less than 1% of the DNS resolvers enforce DNSSEC (cryptographic) validation [13].

We believe that the main reason for the underutilisation of cryptography, is an *illusion of security* against network-based attacks, due to *two false beliefs*. The first false belief, is that in reality, attackers can rarely obtain MitM capabilities, and even when they can, they are reluctant to do so since such activities may lead to detection. We claim that this is incorrect; there are common scenarios where attackers may obtain MitM capabilities, e.g., by accessing wireless communication, by manipulations of the (largely unprotected) routing mechanisms, or by controlling some intermediate device; furthermore, such attacks are often carried out, without detection and repercussions, e.g., BGP route hijacking happens frequently [3].

However, in this article, we focus on the second false belief, which is that current, non-cryptographic, Internet protocols already provide sufficient protection against *typical, common* attackers, and in particular, against *off-path attackers*. Unlike a MitM attacker, an off-path attacker cannot observe or modify legitimate packets sent between other parties, however, it can transmit packets with a spoofed (fake) source IP address - impersonating some legitimate party, as illustrated by *Off-path Oscar* in Figure 1.

Our main goal in this paper is to convince that this second belief is (also) false, and that *current Internet protocols are often vulnerable even to an off-path attacker*. Specifically, we discuss few recent results, that allow off-path attacks on basic Internet protocols: traffic injection into TCP connections and DNS cache poisoning.

<sup>\*</sup>mail@yossigilad.com

<sup>†</sup>amir.herzberg@gmail.com

<sup>‡</sup>haya.shulman@gmail.com

Note that spoofed packets are used in many attacks, most notably, in Denial of Service (DoS) attacks. Significant efforts are made to make spoofing less readily available to attackers, most notably ingress filtering (RFC 3704). However, IP spoofing is still possible via many ISPs, see [1, 6].

The key, to the off-path attacks that we discuss, is *circumvention of challenge-response defenses*, which are often relied upon to distinguish between (spoofed) packets from an off-path attacker and (legitimate) packets from legitimate communication end point. In order to authenticate a response from a server, a client sends a random challenge with the request, which is returned with the response. Since an off-path attacker, which we dub Oscar, cannot eavesdrop on packets exchanged between the server and the client, it appears that he would have to guess the challenge; hence, the (sufficiently long, random) challenge allows to prevent Oscar from crafting a packet with a valid challenge. The security of most Internet applications, e.g., email, web surfing, and most peer-to-peer applications, relies on challenge-response mechanisms, mostly as part of the underlying TCP and DNS protocols. For example, the widely-used web-security mechanisms based on cookies and on the ‘same origin policy’ in general, depend on the security of both TCP and DNS.

Note that, trivially, *challenge-response mechanisms are ineffective against MitM attackers*, since they are able to eavesdrop on the challenges. The false sense of security is due to the two false beliefs mentioned above: that MitM capabilities are ‘rarely practical’ and that existing challenge-response mechanisms, in particular, in TCP and DNS, provide sufficient defenses against the (weaker and common) off-path attackers. The belief that off-path attackers cannot *inject* traffic into a TCP connection is even stated in RFCs and standards, e.g., RFC 4953, discussing TCP spoofing attacks. The reasoning is that TCP specifications and implementations were enhanced to provide security against such adversaries, who are incapable of eavesdropping to communication: modern TCP implementations randomise the 32-bit sequence number [12], and many also randomise the 16-bit client port [19]. To successfully inject data into a TCP stream, the attacker must provide valid values in both fields. Indeed, since its early days, most Internet traffic is carried over TCP - and is not cryptographically protected, in spite of warnings [4, 5, 20].

In contrary to this second belief, we argue that *there are (common) situations allowing off-path attacks*, i.e., where Oscar can circumvent challenge-response mechanisms. One approach [10, 14] exploits *IP fragmentation*, i.e., division of IP packets into several fragments where only one fragment contains the response-field; this allows off-path attacks that block, modify and even intercept fragments or whole packets in some scenarios.

However, in this article we focus on another approach: in order to circumvent deployed challenge-response mechanisms in TCP and DNS, we exploit the fact that to ease deployment, challenge-response defenses for these protocols mostly or wholly *reuse existing fields*; i.e., challenges are fields which already exist in requests and are echoed in responses for some other purpose. In TCP, the ‘challenge’ fields are only reused fields: sequence number and client port; in DNS, there is a short ‘dedicated’ challenge field (16 bits), but since it is insufficient to foil attacks, other fields are ‘reused’ as challenges, including source port, source/destination IP addresses, and the query itself, see RFC 5452.

The use of a randomised (16-bit) source port field, that maps responses to the client process which issued the request, is a widely used ‘best practice’ against off-path attacks; see RFCs 6056, 5452 and [16].

We argue that such *dual-use* of an existing field for challenge-response, while conveniently allowing deployment of defenses only on the client side without requiring coordinated adoption by the server, is often vulnerable. Specifically, we discuss attacks that allow an off-path attacker to *learn the source port* and other ‘dual-use’ fields. This allows off-path TCP injection and DNS cache poisoning.

## 1.1 History of Off-Path Attacks

TCP and DNS are basic protocols, and off-path attacks on their authenticity - TCP injection and DNS poisoning - impact almost all Internet applications. As such, it is a common belief that they ensure integrity against an off-path attacker. However, security against off-path (or MitM) attackers was not of the original design goals of these protocols, and only minimal changes were done to the specifications to support challenge-response defenses, e.g., selecting identifiers at random.

Indeed, over the years, significant attention and efforts were dedicated to validating and improving the off-path security of TCP and DNS - and numerous off-path attacks were launched, some of them widely publicised. In Figure 2 we present a ‘time-line’ of important attacks and security improvements, for both TCP (upper row) and DNS (lower row).

The time-line begins in 1985 with Morris publication of TCP injection attack based on the use of predictable sequence numbers [20], and Bellovin’s seminal paper from 1989 [4], pointing out that security should *not* be based on the presumed off-path protection of DNS and TCP. Bellovin presented vulnerabilities of (some) TCP implementations to off-path attacks, and discussed potential exploits and defenses.

Unfortunately, in spite of these warnings, until 1995 most TCP stacks still used trivially-predictable *initial sequence number (ISN)*. This changed only after the infamous attack by Mitnick on Shimomura [23], that utilized a TCP injection. After the attack, many implementations changed to ‘less predictable’ ISN choices. However, in 2001, Zalewski showed that most implementations are still ‘sufficiently predictable’, allowing off-path attacks; this motivated adoption of more random choice of ISNs in most operating systems, as standardised in RFC 6528.

One additional TCP injection attack was presented in 2004 by Watson [24]. This attack only injected a ‘RST’ packet, breaking up a connection, and focused on long-lived connections using known client (and server) ports and addresses, as used at the time, in particular, by the Internet routing protocol BGP. To address this concern, many TCP implementations began also using ‘unpredictable’ client ports.

In 2007, there were two surprising results: (1) a TCP injection attack presented by the pseudonym author *klm* in Phrack magazine [18], and (2) a DNS poisoning attack exploiting poor random-number generators [17]. Both attacks were clever and significant although with limited scope. In particular, the attack on TCP worked only against Windows machines, connected directly to the Internet (rather than via firewall, as usually is the case), and did not handle concurrent connections.




Figure 2: Time-line of DNS Poisoning and TCP Injection attacks.

Kaminsky presented an even more significant DNS poisoning attack in 2008, which allowed efficient off-path poisoning of most DNS resolver implementations at the time [16] (see Section 2). The response to this attack was rapid adoption of additional ‘patches’, mostly, more challenge-response fields, increasing the randomisation and therefore (hopefully) making the attack impractical; the most notable patch was *source port randomisation (SPR)*; see RFC 5452.

Following 2008, there were several years without additional off-path attacks; Kaminsky’s attack was addressed by SPR and other ‘patches’ to resolvers, and klm’s attack was impractical and not widely known (due to the unusual venue). This changed dramatically in 2011-2013, with the publication of *eight* new off-path attacks. The first, in 2011, was an attack on fragmented IP traffic [10]. This was followed, by *two* new DNS poisoning attacks (in 2012 [15] and 2013 [14]), a connection-exposing attack [9] and *four* TCP injection attacks [8, 11, 21, 22]. We discuss some of these attacks in this paper.

## 1.2 Malicious Agents

Some of the off-path attacks require, in addition to the spoofing ability, also a *malicious agent* in the victim’s network or host. We now explain the different agent models.

A *user-level zombie* is a machine controlled by the adversary, e.g., compromised by malware, in the victim’s network. There appears to be a significant amount of attacker controlled hosts (zombies) on the Internet.

A *puppet* [2] is weaker agent: a restricted malicious script or applet running in web-browser sandbox. Attacks relying on a puppet agent require (only) that a client in the victim network ‘surfs’ to the attacker’s web-site, enabling the adversary to run such a script. The script is restricted by *same origin policy* (described in RFC 6454), and can only communicate via the browser, i.e., request (and receive) HTTP objects (no access to TCP/IP packet headers).

Puppets are usually easier to obtain and control compared to *zombies*, since browsers normally run scripts automatically upon opening a web-site, while zombies require installation (of malware).

## 2. DNS CACHE POISONING

Challenge-response mechanisms allow DNS resolvers to authenticate legitimate DNS responses, thus preventing cache poisoning by off-path attackers. Until the year 2008 the only (widely deployed) challenge-response mechanism was the 16-bit *transaction identifier (TXID)*.

### 2.1 Kaminsky’s DNS Cache Poisoning

In 2008, Kaminsky [16] presented an efficient cache poisoning attack against resolvers which authenticated responses

using only the TXID. We briefly describe this attack. The attack assumes a known source port, and for concreteness we used source port 53. The steps of the attack, illustrated in Figure 3, are the following:

- (1) the attacker triggers a DNS request for a random sub-domain of the victim domain **1\$.foo.com**. There are different techniques to trigger DNS requests, e.g., a benign client visits a web page of the attacker; the web page contains an object from a victim domain, which triggers a DNS request.
- (2) the recursive DNS resolver receives a DNS request and sends a DNS request to the target name server.
- (3) the attacker then sends  $2^{16}$  responses with spoofed source IP (of the name server); each response is a **referral**, mapping the name server **ns.foo.com** to 6.6.6.6, an IP address controlled by the attacker.
- (4) the response containing the correct TXID is accepted, cached and sent to the client.
- (5) authentic DNS response is ignored, since there is no matching pending request.

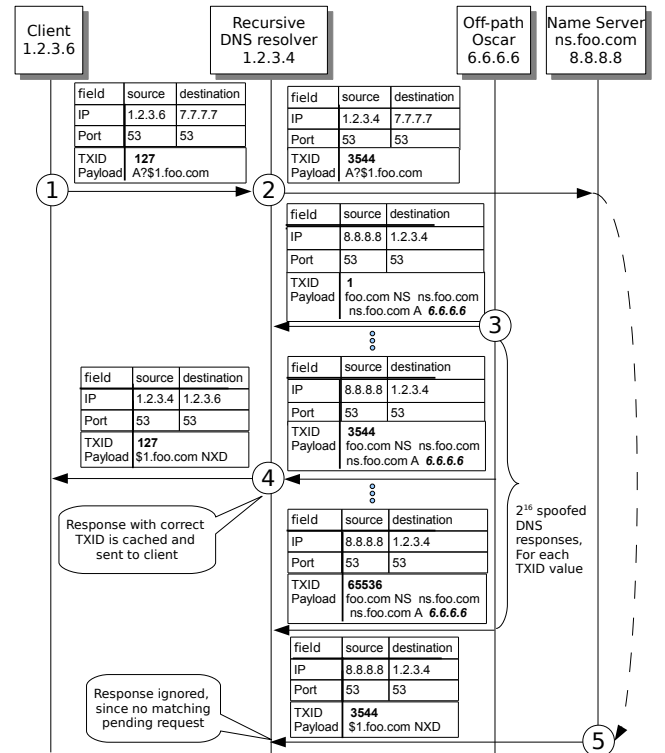


Figure 3: Kaminsky’s Attack.

Following to Kaminsky’s attack, additional challenge-response mechanisms were proposed to increase the entropy in DNS requests (RFC 5452). The most popular mechanism, supported by most resolvers, is source port randomisation (SPR). SPR in tandem with TXID, produce a search space of  $2^{32}$  and were believed to provide sufficient protection against poisoning by off-path adversaries, since this requires the attacker to not only guess the TXID but also the source port correctly. This reduced the motivation for deployment of the cryptographic defense against poisoning, DNSSEC (RFCs 4033-4035). We found different techniques, [14, 15], to circumvent the popular defenses against poisoning by off-path attackers. In this work we show a simple technique from [15], which applies to common network scenarios, where the DNS resolvers are located behind a NAT device (as in Figure 1), and uses side-channels for ports’ prediction.

## 2.2 Vulnerability of Resolvers Behind NAT

Network Address Translation (NAT) devices that assign predictable ports to outgoing packets can expose resolvers to cache poisoning attack. We show that even NAT devices that sufficiently randomise outgoing ports may expose resolvers to attacks.

### Per-Destination NAT.

NAT devices modify ports in outgoing packets, and maintain mappings between internal and external ports.

We focus on the common *per-destination* NAT; see [15] for other NAT devices. A per-destination NAT operates as follows: for a tuple defined by  $\langle \text{src-IP}:\text{src-port}, \text{dst-IP}:\text{dst-port}, \text{protocol} \rangle$ , it selects the first port at random, and subsequent ports are increased sequentially (for that tuple).

### Predict-then-Poison Attack.

The off-path attacker, Oscar, controls a *zombie*, i.e., non-privileged malware, that runs on a client host in the LAN.

The first phase of the attack is *port prediction* during which the zombie and Oscar collaborate to expose the port that will be allocated to the DNS request of the DNS resolver. Then during the second phase, Kaminsky’s attack is launched. We present only the *predict* phase of the attack, which steps are illustrated in Figure 4; the steps of Kaminsky’s cache poisoning attack are in Figure 3.

(1) The zombie sends a packet to create a mapping in the NAT table; in the example in Figure 4 we assume arbitrarily that port 6666 was selected.

(2) Then, Oscar at address 6.6.6.6 sends  $2^{16}$  packets with a spoofed source IP of 8.8.8.8, s.t., each is sent to a different port of the NAT and each contains the destination port in payload; only the packet with port 6666 arrives at the zombie.

(3) The zombie increments this port by 1, in our example the result is 6667, and sends it to Oscar; this is the external port that will be allocated by the NAT to the subsequent DNS request of the resolver to the victim name server.

This phase allows to bypass the SPR defense.

## 3. TCP INJECTIONS

The recent off-path TCP injection attacks operate in two phases:

◊ **Learn Connection 4-Tuple.** Oscar, the off-path attacker,

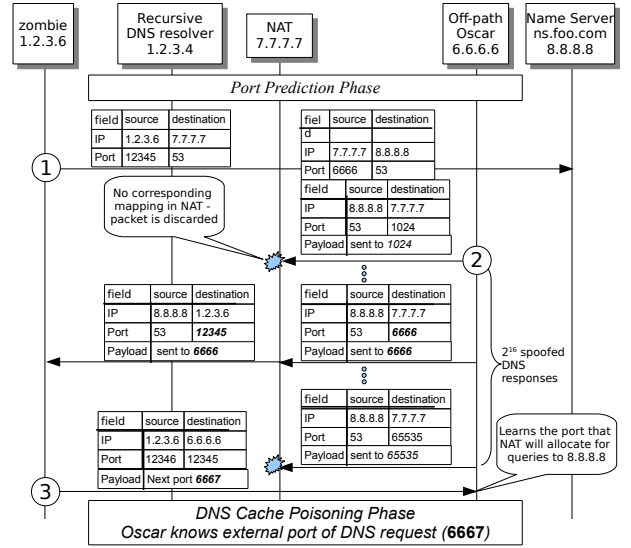


Figure 4: Predict-then-Poison DNS Attack Assuming Per-Destination NAT.

learns the four parameters of a TCP connection between a client and a server, that is, their respective IP addresses and ports.

◊ **Learn Sequence Number(s).** Oscar learns the current sequence number, for packets sent from the server to the client. In some attacks, Oscar also learns the sequence number for packets from the client to the server. In this phase Oscar bypasses the initial sequence number randomisation defense, which is implemented in all modern TCP/IP stacks.

Table 1 surveys the techniques used in recent injection attacks for the two phases and their requirements (in parentheses). In the remainder of this section we present a simple implementation for each phase.

### 3.1 Learn Connection 4-Tuple

In order to launch an injection attack, Oscar must first identify a TCP connection between the victim client and server.

We describe the method of [8], which uses a puppet (script restricted by browser sandbox) running on the client machine to *open* such a connection. Since Oscar chooses the server, the server’s IP address and port are known. To find the client’s IP, the puppet sends a request to Oscar’s site; this request contains the client’s IP address.

The final challenge of this phase is to detect the client port. Many clients, in particular, those running Windows, assign ports to connections *sequentially*. The attack of [8] uses the puppet to open a connection to Oscar’s remote site before and after opening the connection to the victim server; sequential port assignment allows Oscar to learn the client’s port: Oscar observes  $p_1$  and  $p_2$ , the client ports used in the connection to his sites. If  $p_2 = p_1 + 2$ , then he identifies that the connection to the server is via port  $p_1 + 1$  (otherwise Oscar restarts the attack).

### 3.2 Learn Sequence Numbers

The next step after identifying the victim-connection, is learning one or both connection’s sequence numbers. Observing the sequence numbers directly from traffic requires

	Learn Connection 4-tuple	Learn Sequence Numbers
klm [18]	Active probing for connection (Windows client, no firewall)	Side channel (Windows client)
Qian et al. [21, 22]	Monitor connections, e.g., with <i>netstat</i> (Malware)	Read client system-counters, (Malware; in [21] also seq. # checking firewall)
Gilad and Herzberg [8]	Establish connection, exploit sequential port allocation impl. (Puppet, Windows client)	Side channel (Puppet, Windows client)
Gilad and Herzberg [11]	Establish connection, client port de-randomization (Puppet, client behind firewall)	Exploit browser behavior, (Puppet, no TLS/SSL)

**Table 1: Off-Path TCP Injection Attacks: Building Blocks. In parentheses: requirements.**

an on-path attacker (eavesdropping capability); off-path attacks use different methods to infer the sequence numbers. We focus on the technique of [11] which exploits an *under-specification of HTTP* to learn the client’s sequence number.

**BACKGROUND.** As of HTTP 1.1, clients can send multiple requests to the same web-server in *pipeline* over a single (‘persistent’) HTTP connection. In order to allow browsers to match between each response and the corresponding request, the server sends the responses exactly in the order in which it received the requests. The browser (client) keeps a FIFO queue of pending HTTP requests for each connection, and handles them one by one, as follows. To handle a request, the browser reads the bytes in the TCP connection’s receive-buffer (when they become available). The browser expects to find the matching response in the beginning of TCP’s receive-buffer and parses the response.

The HTTP standard does not specify what the browser should do when the receive-buffer contains data which is *not* a valid (‘parsable’) HTTP response. Many modern browsers (including Chrome, Firefox and Internet Explorer) handle this situation as follows: the browsers treat *all available data* in the receive-buffer as payload of a response with the following ‘default’ HTTP header:

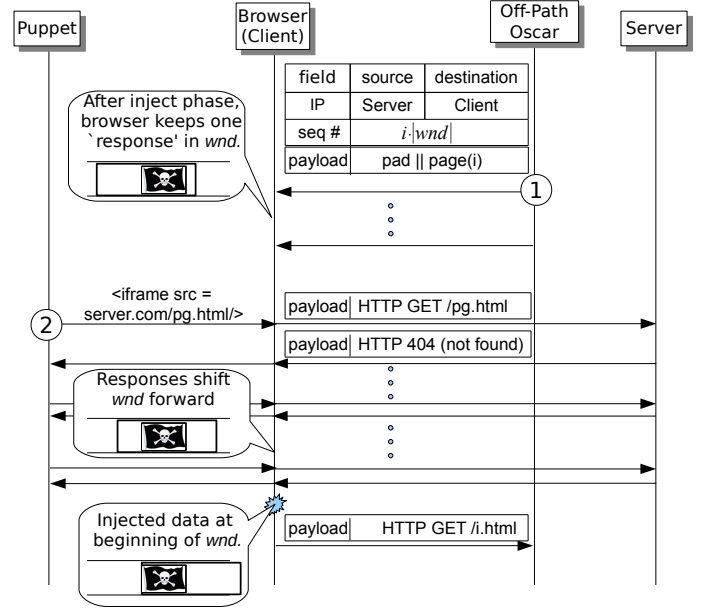
```
HTTP/1.1 200 OK
Content-Type: text/html; charset=us-ascii
Content-Length: available-data-size
```

The browsers return this ‘response’ to the requesting module, normally, the rendering engine or a script/applet.

**SEQUENCE NUMBER LEARNING.** The learning phase has two steps: *Inject* and *Observe*, see Figure 5. In the inject step, Oscar injects data into the stream of HTTP responses that the server sends to the client. This data is read in the observe step, which allows Oscar to determine the server’s sequence number.

(1) **Inject step.** Let  $wnd$  denote the browser’s receive-buffer for the connection and  $|wnd|$  denote its size. In order to inject the data, Oscar sends to the browser  $\frac{2^{32}}{|wnd|}$  packets, spoofed to appear to be from the server (on its victim-connection with the client). The  $i^{th}$  packet has server sequence number  $i \cdot |wnd|$ ; since the sequence number field is 32-bits long, exactly *one* of these packets has a ‘valid’ sequence number, which falls within  $wnd$ ; all the other packets are discarded by the client. Each of Oscar’s packets contains as payload  $page(i)$  which is a simple web-page defined as follows:

```
<HTML><BODY>
<iframe src = "oscar.com/i.html" />
</BODY></HTML>
```



**Figure 5: Sequence Number Learning Technique.**

(2) **Observe step.** In this step, the puppet makes prevalent requests to the server, until it reaches the data injected by Oscar in the previous step. Each server-response that arrives at the client shifts  $wnd$  forward; after several such responses arrive, there is no gap of unreceived bytes between the injected data and the beginning of  $wnd$ . Then, the browser reads the injected-response, assuming that it corresponds to the request.

The last server-response usually overwrites the beginning of the injected data, therefore, the injected ‘response’ will usually be corrupt. However, as noted above, many browsers handle the injected data as payload ‘wrapped’ with a default header. When the browser renders  $page(i)$ , it tries to retrieve  $i.html$  from Oscar’s web-site (see Figure 5); providing to Oscar the value of  $i$ . In order to keep  $page(i)$  intact despite the overwrite, when Oscar sends  $page(i)$  in the inject step, he prepends to it an easily removable pad. The value of  $i$  allows Oscar to compute the next server sequence number that the client expects.

## 4. EXPLOITING INJECTION/POISONING

To conclude our discussion of off-path TCP injection and DNS poisoning attacks, we briefly discuss some potential exploits.



Exploiting *DNS poisoning* is straightforward, e.g., see [16]. Both users and programs use DNS extensively to resolve domain names - to obtain the IP address and communicate with a server, or to obtain other DNS-indexed resource, often security related, e.g., SPF policies or blacklists. DNS poisoning allows circumvention of these mechanisms, e.g., ‘hijacking’ of connection requests meant for legitimate servers. In particular, ‘hijacking’ can allow *phishing*, where a user thinks he interacts with a trusted site, while he actually deals with a fake site (exposing credentials, installing malware, etc.). Furthermore, the malicious mapping is kept for the time (TTL) specified in the record, and in the common case of a single cache used by many clients, the poisoning impacts all of them.

Exploiting *TCP injections* is more challenging, since TCP is a transport protocol and does not involve caching. However, in common scenarios, TCP injections can allow critical exploits. In particular, TCP injections suffice to *circumvent Same Origin Policy*, hijack ‘cookies’ and cause execution of malicious scripts (XSS). Furthermore, to cause long-term impact similar to DNS poisoning, attackers can exploit caching of objects by web caches. By crafting the HTTP headers of his response, Oscar can cache the spoofed page and specify that it will be out-dated only after a long time; furthermore, when using a web-cache, the impact can be over many users. See details in [11].

## 5. CONCLUSIONS

The vulnerabilities presented in this work show that relying on challenge-response mechanisms against off-path attackers can often be circumvented. Specifically, the techniques discussed in this work allow off-path attackers to circumvent the main challenge-response defenses: *source port randomisation* and *initial sequence number randomisation*.

Our message is that defenses should be designed and analysed against the strongest MitM attackers. In particular, to prevent these, and other unforeseen, vulnerabilities we recommend deployment of systematic, cryptographic defenses for TCP and DNS.

## 6. REFERENCES

- [1] Advanced Network Architecture Group. ANA Spoofer Project. <http://spoofer.csail.mit.edu/summary.php>, 2012.
- [2] S. Antonatos, P. Akritidis, V. the Lam, and K. G. Anagnostakis. Puppetnets: Misusing Web Browsers as a Distributed Attack Infrastructure. *ACM Transactions on Information and System Security*, 12(2), 2008.
- [3] H. Ballani, P. Francis, and X. Zhang. A Study of Prefix Hijacking and Interception in the Internet. In *ACM SIGCOMM Computer Communication Review*, volume 37, pages 265–276. ACM, 2007.
- [4] S. M. Bellovin. Security Problems in the TCP/IP Protocol Suite. *Computer Communication Review*, 19(2):32–48, apr 1989.
- [5] S. M. Bellovin. A Look Back at “Security Problems in the TCP/IP Protocol Suite”. In *ACSAC*, pages 229–249. IEEE Computer Society, 2004.
- [6] R. Beverly, A. Berger, Y. Hyun, and K. C. Claffy. Understanding the Efficacy of Deployed Internet Source Address Validation Filtering. In A. Feldmann and L. Mathy, editors, *Internet Measurement Conference*, pages 356–369. ACM, 2009.
- [7] CAIDA. Anonymized Internet Traces 2012 Dataset. [http://www.caida.org/data/passive/passive\\_2012\\_dataset.xml](http://www.caida.org/data/passive/passive_2012_dataset.xml), 2012.
- [8] Y. Gilad and A. Herzberg. Off-Path Attacking the Web. In *USENIX Workshop on Offensive Technologies*, pages 41 – 52, 2012.
- [9] Y. Gilad and A. Herzberg. Spying in the Dark: TCP and Tor Traffic Analysis. In S. Fischer-Hübner and M. Wright, editors, *Privacy Enhancing Technologies Symposium*, volume 7384 of *Lecture Notes in Computer Science*, pages 100–119. Springer, 2012.
- [10] Y. Gilad and A. Herzberg. Fragmentation Considered Vulnerable. *ACM Transactions on Information and System Security (TISSEC)*, 14(4):16:1–16:31, April 2013. A preliminary version appeared in WOOT 2011.
- [11] Y. Gilad and A. Herzberg. When Tolerance Becomes Weakness: The Case of Injection-Friendly Browsers. In *Proceedings of the International World Wide Web Conference*, May 2013.
- [12] F. Gont and S. Bellovin. Defending against Sequence Number Attacks. RFC 6528 (Proposed Standard), Feb. 2012.
- [13] O. Gudmundsson and S. D. Crocker. Observing DNSSEC Validation in the Wild. In *SATIN*, March 2011.
- [14] A. Herzberg and H. Shulman. Fragmentation Considered Poisonous. *CoRR*, abs/1205.4011, 2012.
- [15] A. Herzberg and H. Shulman. Security of Patched DNS. In S. Foresti, M. Yung, and F. Martinelli, editors, *ESORICS*, volume 7459 of *Lecture Notes in Computer Science*, pages 271–288. Springer, 2012.
- [16] D. Kaminsky. It’s the End of the Cache As We Know It. In *Black Hat conference*, August 2008. <http://www.blackhat.com/presentations/bh-jp-08/bh-jp-08-Kaminsky/BlackHat-Japan-08-Kaminsky-DNS08-BlackOps.pdf>.
- [17] A. Klein. OpenBSD DNS Cache Poisoning and Multiple O/S Predictable IP ID Vulnerability. [http://www.openbsd.support.com.ar/books/OpenBSD\\_DNS\\_Cache\\_Poisoning\\_and\\_Multiple\\_OS\\_Predictable\\_IP\\_ID\\_Vulnerability.pdf](http://www.openbsd.support.com.ar/books/OpenBSD_DNS_Cache_Poisoning_and_Multiple_OS_Predictable_IP_ID_Vulnerability.pdf), October–November 2007.
- [18] klm. Remote Blind TCP/IP Spoofing. Phrack magazine, <http://www.phrack.org/issues.html?id=15&issue=64>, 2007.
- [19] M. Larsen and F. Gont. Recommendations for Transport-Protocol Port Randomization. RFC 6056 (Best Current Practice), Jan. 2011.
- [20] R. T. Morris. A Weakness in the 4.2BSD Unix TCP/IP Software. Technical report, AT&T Bell Laboratories, Feb. 1985.
- [21] Z. Qian and Z. M. Mao. Off-Path TCP Sequence Number Inference Attack. In *IEEE Symposium on Security and Privacy*, pages 347–361, 2012.
- [22] Z. Qian, Z. M. Mao, and Y. Xie. Collaborative TCP Sequence Number Inference Attack: How to Crack Sequence Number Under a Second. In *Proceedings of the ACM Conference on Computer and*

*Communications Security*, pages 593–604, New York, NY, USA, 2012. ACM.

- [23] T. Shimomura and J. Markoff. *Takedown: The Pursuit and Capture of Kevin Mitnick, America's Most Wanted Computer Outlaws - by the Man Who Did It*. Hyperion Press, 1st edition, 1995.
- [24] P. Watson. Slipping in the Window: TCP Reset Attacks. presented at CanSecWest, <http://bandwidthco.com/whitepapers/netforensics/tcpip/TCP%20Reset%20Attacks.pdf>, October 2004.